The Return to Capital for Small Retailers in Kenya:

Evidence from Inventories*

Michael Kremer, Jean N. Lee, Jonathan Robinson, and Olga Rostapshova

July 2011

PRELIMINARY AND VERY INCOMPLETE: PLEASE DO NOT CIRCULATE OR CITE WITHOUT PERMISSION

Abstract

Standard textbook models suggest risk-adjusted rates of return should be equalized across activities within firms, and across firms. However, testing this implication is typically complicated by the difficulty in measuring rates of return. In this paper, we take advantage of the characteristics of the retail industry in rural Kenya to create estimates and bounds on the rate of return to inventories in a set of retail firms. Using administrative data on whether firms purchased enough to take advantage of quantity discounts from wholesalers, we estimate a lower bound on rates of return for the median shop of well over 100 percent per year. A second approach measures the return to increased investment by surveying shops on a regular basis about "stockouts" (lost sales due to insufficient inventory). Though average returns are much lower from this approach (likely because firms fear losing customers due to lost goodwill), we are able to reject the hypothesis that the marginal rates of return are equal across shops. We examine alternative explanations of why we observe such high and heterogeneous marginal returns to capital and find evidence that credit constraints are not the only explanation. Though preliminary, our results suggest significant departures from the notion of firms as purely rational profit makers.

* Kremer: Harvard University, Brookings Institution, Center for Global Development (CGD), J-PAL, NBER, and BREAD; Lee: Harvard University; Robinson: UCSC, J-PAL, and BREAD; Rostapshova: Harvard Kennedy School. We thank David Autor, Austan Goolsbee, David McKenzie, Sendhil Mullainathan, Kevin M. Murphy, Ben Olken, Kim Siegal, Bent Sorensen, John Sutton, seminar participants at CEPR, NBER Africa, RAND, Stanford, UCB ARE, UCSC, Yale, the 2011 BREAD summer workshop and the World Bank Microeconomics of Growth conference for helpful comments. Kenzo Asahi, Pia Basurto, Dan Bjorkegren, Conner Brannen, Elliott Collins, Sefira Fialkoff, Katie Hubner, Eva Kaplan, Anthony Keats, Jamie McCasland, and Russell Weinstein provided excellent research assistance. We gratefully acknowledge funding from the SEVEN Foundation and the NBER Africa program. We thank IPA-Kenya for administrative support.

1. Introduction

Standard textbook economic models suggest that the risk-adjusted rate of return should be equalized across activities within a firm. If capital markets function well, rates of return should also be equalized across firms, both within and even across countries. While it is clear that various frictions interfere with perfect equalization of rates of return across firms, it is not clear how big the departures from this benchmark are, and which departures are most important.

In addition, it is often difficult or impossible to directly measure rates of return to capital, particularly at the margin. In this paper, we take advantage of the structure of the retail industry among a subset of Kenya retailers to measure the rate of return to inventory capital. We are able to identify specific investment opportunities that are available to retail firms and directly compute the return that could be realized from these investments. The data imply very high marginal rates of return on average, and provide evidence for economically and statistically significant heterogeneity in marginal rates of return across shops.

In our first empirical strategy, we calculate the bounds on the rate of return to investment in inventory for over 2,000 shops in Western Kenya using a complete database of shop purchases from a major distributor of retail goods. We infer bounds on the rates of return to investments in inventory that could be achieved if shops shifted the timing of their purchases to take advantage of quantity discounts offered by the distributor. If shops have investment opportunities that exhibit diminishing returns at least locally, the average return on these incremental investments will be a lower bound on the marginal rate of return. Preliminary estimates using this approach suggest that approximately 40% of firms have unexploited investment opportunities that would yield a real rate of return over 100 percent on an annualized basis.

In our second empirical approach, we directly measure the expected rate of return to an incremental investment in inventory for small retail firms in Western Kenya. We collected detailed

panel data on inventory decisions, sales, and stockouts (lost sales in which a customer asks for a product that it out of stock and does not accept a replacement) for a sample of 142 small rural retail firms in 16 towns in Western Kenya. By measuring daily stockouts over a period of several months, we are able to measure the probability that an additional unit of inventory would have been sold in a given time period, had the shopowner bought it at the beginning of the period. In this way, we are able to estimate marginal rates of return to inventory investment by calculating the expected marginal benefit from holding an additional unit of inventory (the markup multiplied by the probability that the marginal unit would sell during the relevant time period), and comparing this to the marginal cost of obtaining an additional unit (the wholesale price multiplied by the cost of financing).

We focus our analysis on cell phone top-up cards, for several reasons. First, phone cards have fixed wholesale and retail prices and negligible storage and depreciation costs, and are not substitutable across brands. Second, phone cards are kept behind the counter in the shops we survey, so lost sales can be measured and theft is rare. However, one problem with this approach is that the costs of stocking out of phone cards might be quite high – shops may lose customer goodwill, and some evidence suggests that phone cards are "leader" items which get people into shops. Likely for this reason, we find modest returns to holding additional phone cards: the average shop in our sample would achieve a real rate of return of 33 percent to a marginal increase in inventory per year, with 17.6 percent of shops having returns greater than 50%. If lost customer goodwill or other sales of complementary goods are significant, this will be a lower bound on the rate of return.

If one treats the calculated rates of return as estimates rather than bounds, or assumes that all these bounds are equally tight because the cost of lost goodwill and other sales is similar across shops, we can then test whether these marginal rates of return are equal across shops, and estimate

the degree of heterogeneity in rates of return under some assumptions about the underlying distribution of rates of return. Using a variety of tests, we reject the hypothesis of equalization of marginal rates of return across shops, suggesting some misallocation of capital in these markets. In the stockout analysis, we find evidence that the standard deviation of the population distribution of annual rates of return may be as high as 98 percent.

We explore the extent to which the estimated rates of return (from both methods) may reflect high rates of return to capital or behavioral anomalies in several ways. First, we examine how the returns vary with background characteristics, and we find little evidence of a correlation with characteristics which are likely to be correlated with credit constraints (for example, asset or land ownership). Second, we examine whether returns from the two methods are correlated, and find a weak relationship. Relatedly, we also separately estimate rates of return (from the stockout analysis) to two different brands of phone cards, Celtel and Safaricom, for each shop in the sample. Again, the cross-brand correlation is weak. Third, we find evidence that firms are less likely to report stockouts the longer they have been followed in our monitoring surveys. An explanation for this finding is that our monitoring visits (in which we asked firms about their lost sales from stockouts) might have made stockouts more salient to shopkeepers and they might have adjusted their behavior in response. However, this finding is speculative in that the experiment was not designed for this purpose and so variation in exposure to the survey is non-random..

This paper contributes to a growing empirical literature on marginal rates of return to capital. Lucas (1990) famously noted that the simplest calibration exercise assuming a common aggregate production function suggests that the marginal rates of return to capital must differ dramatically between the rich and poor countries of the world. A recent paper by Caselli and Feyrer (2007) argues that the aggregate country level data on capital share of income, output, capital stock, are consistent with equalization of financial marginal rates of return across countries, after accounting

for payments to previously unobserved factors (such as land and natural resources) and differences in prices of investment goods across countries.

The development literature tends to find evidence for high and variable marginal rates of return to capital. The approaches in this literature vary, but in general the estimated returns to capital are quite high, well above typical estimates for the developed world. First, several early studies (for example, Aleem 1990) note that marginal rates of return must exceed the high interest rates at which people and businesses are willing to borrow, but may include some borrowing to smooth consumption as well as borrowing for productive investments .

A number of recent studies use detailed microeconomic data and find high rates of return. Several notable recent studies measure marginal returns by experimentally increasing capital stocks through cash or capital grants. These studies find returns on the order of 5-20% per month in Sri Lanka (de Mel, McKenzie, and Woodruff 2008), Mexico (McKenzie and Woodruff, 2008), and Ghana (Fafchamps et al. 2011), respectively. McKenzie and Woodruff (2006) similarly find large returns non-experimentally to small entrepreneurs. Banerjee and Duflo (2008) examines policy shocks to directed lending to much larger firms in India and concludes that marginal rates of return to capital exceed 70 percent for those firms affected by the changes. Exploiting county-level variation in credit supply due to the Community Reinvestment Act, Zinman (2002) estimates gross rates of return to capital in the US on the order of 20-58 percent per year. In a study which uses a somewhat similar approach to ours, Anagol and Udry (2006) take the elegant approach of using data on prices of used car parts of varying expected lifetimes to infer a 60 percent annual discount rate for taxi drivers in Ghana. The authors also find an annual rate of return of 150 to 250 percent to pineapple cultivation in Ghana.

Similarly, this paper finds high returns using a different methodology. Our paper also finds evidence of significant heterogeneity in returns within a very narrowly defined industry, and

potential capital allocation. Our work is thus related to papers such as Hsieh and Klenow (2009), who find significant misallocation of capital among manufacturing firms in China and India.

Though our results remain quite preliminary, our data should eventually enable us to shed some light on *why* returns differ so drastically. The two most likely explanations are that (1) some firms are credit constrained, and (2) that firms fail to take advantage of profitable opportunities for behavioral reasons. Though our current evidence is quite speculative, our data suggests that credit constraints are unlikely to provide a full explanation for the pattern of findings. We are in the process of conducting a field experiment to examine whether simply telling firms about the lost returns changes behavior. Though we do not have results from that intervention, we do have some evidence from debriefing interviews which suggests that lack of attention and/or behavioral issues may explain much of what we find. These results are discussed at the end of the paper.

This paper is organized as follows. Section 2 discusses the context of the small-scale retail sector in Kenya. Section 3 describes the distributor data and the estimated rates of return implied by the bulk discount method. Section 4 introduces the stockout survey and data, presents the rates of return implied by the stockout data, suggests a framework for interpretation and shows that these data imply very high marginal rates of return for a nontrivial fraction of firms. Section 5 describes survey data on demographic, personality and entrepreneurial characteristics of the shop owners and describes their relationship to the estimated rates of return. Section 6 concludes.

2. The Small-Scale Retail Sector in Kenya

The small-scale retail sector comprises a significant share of economic activity in Kenya, particularly in rural areas. Daniels and Mead (1998) estimate that small and medium enterprises with 10 or fewer employees (not including agriculture and mineral extraction industries) comprise 12-14% of total Kenyan GDP, and that a quarter of this contribution comes from the retail trade.

We focus on a category of retail shop in Western Kenya called *dukas* in Kiswahili, which are small scale enterprises that are typically owner-operated, often by women and those with some secondary education. These shops are ubiquitous in market centers and small towns in the region, and are often located in clusters: adjacent to or in close proximity to several competing shops. They typically sell a relatively homogeneous set of household products such as perishable and non-perishable foodstuffs, soaps, detergents, cooking fat, sodas, phone cards, and other household items. This is a poor area, and people consume a pretty small set of goods, so there is almost no variation across shops in the types of products they carry. The decision about how much inventory to buy is the most important optimization problem that most of these shopowners face. Most of these goods have low depreciation and are non-perishable, and the storage costs are often not substantial, especially on the margin.

Products are kept behind a counter (and often behind a set of metal bars) and all transactions and transfers of goods are mediated through the store operator. This means that we potentially have information on stockouts, although people may see certain goods out of stock and not ask for them. Phone cards, however, are kept below the counter so that customers are unable to know that they are out of stock without inquiring with the shopkeeper. Operators deal with a number of suppliers for the different goods they sell, but typically do business with a single supplier for each type of good.

Many goods are delivered on a regular schedule. Distributors are based in larger, semi-urban towns, and deliveries are made several times a week, depending on the product. For shops that are located in or near these larger towns, it is possible to restock from the distributor immediately if a stockout occurs or is soon to occur. For this study, however, we selected shops that are located too far from their distributors to make travel for restocking profitable. In addition, shops are able to restock certain products at any time by purchasing from a wholesaler that is located nearby. The

disadvantage of restocking from these wholesalers is that they usually offer a smaller discount from the retail price than do distributors.

In this study, we focus our bulk discount analysis on non-perishable food items and household goods (e.g. vegetable cooking fat, soup mix, soap, and margarine) and our stockout analysis on top-up cards for cellular phone service. These products differ in their typical method of distribution. For the shops in our sample, non-perishable food items and household goods are purchased either from the distributor or from a wholesaler, while phone cards are purchased exclusively from distributors. For our first empirical strategy, we obtained data from 2005-2011 on all purchases from a major distributor that supplies over 100 household goods to retailers within the geographic area. This distributor makes deliveries to each shop once a week and offers the retailers bulk discounts based on total purchase amount at the time of delivery. Before 2010, the discounts range from 0.5% on purchases greater than 5,000 Ksh (approximately $67 – the exchange rate was approximately 75 Ksh to $1 US during this time period) to 1.5% on purchases greater than 10,000 Ksh (about $133), substantial relative to typical retail markup of about 10%. After January 2010, these thresholds changed to 7,500 Ksh, 10,000 Ksh, and 20,000 Ksh respectively.

Our second empirical method uses stockouts of phone cards. Phone cards are a high volume commodity and are carried by many shops. The wholesale and retail prices are fixed and the markup is approximately 4-5%. There are two brands of top-up cards which are specific to the major cellphone carriers which were operating during our data collection period: Celtel and Safaricom. Each brand had several denominations of cards. Celtel cards came in 40, 100, 200, 300, 600, and 1200 Kenyan shilling (Ksh) denominations. During this time period, a small number of shops also had a technology which allowed them to sell cards in arbitrary denominations. Safaricom cards came in 50, 100, 250, 500, and 1000 Ksh denominations. The brands were not substitutable for each other, though there is substitutability across denominations within a brand. Because most consumers

buy the smallest available denomination, there is rarely substitution across denominations in the event in which a shopowner runs out of inventory for the desired denomination. One feature of the phone card distribution system that complicates our analysis is that goods must be purchased in discrete order sizes. For example, cards must often be purchased in packs of ten. Shops buy a minimum of 5 to 20 cards at time, depending on the denomination of the product. For this reason, we calculate the expected profit from holding an additional order of ten cards rather than the return to one marginal card.

3. Bulk Discount Analysis

3.1 Distributor Data

We analyze sales data from a major distributor of retail goods in Western Kenya. These data contain detailed records of purchases between January 1, 2005 and June 30, 2011 for purchases that are less than 100,000 Ksh in value, a rough cutoff which excludes very large wholesalers. During this period, the distributor supplied 160 different household goods. While goods such as eggs, bread, milk and a set of other household goods are distributed separately and not observed in these data, the products in our data (including cooking fat, soap, detergent, soup mix, etc.) make up roughly 30% of a shop's inventory.

We observe the name of the shop, date of the purchase, the quantity purchased of each product, the unit prices, the actual prices paid, the Value Added Tax (VAT) paid, and any discounts received for each purchase. The shop identifiers also include some geographic information. Many shops purchase infrequently from the distributor, close or change names so we do not observe many of these shops for long enough to tightly estimate their returns. In particular, if shops do not reappear near the ends of the dataset we cannot estimate returns for those dates. We therefore restrict analysis to those shops that appear in the data for a sufficient period of time, which we

arbitrarily define as being observed for at least 20 purchases. This will bias the sample towards bigger firms that will likely have lower rates of return than the marginal shop. In addition, there are also some minor issues in matching names in the distributor data, again biasing towards bigger shops that reappear frequently.

The average shop satisfying these inclusion rules makes 63 purchases in the data, and the average length of time between the first and last purchase in the data is 916 days. The average shop in the sample invests about 3,700 Ksh per visit in products sold by this distributor, about 2.2 times a month (although note that the distribution is skewed). Summary statistics for this data appear in Table 1.

Before 2010, shops receive a 0.5 percent discount if their total bill including VAT exceeds 5,000 Ksh, a 1 percent discount if their bill exceeds 7,000 Ksh, and a 1.5 discount if their bill exceeds 10,000 Ksh. These thresholds changed to 7,500, 10,000, and 20,000 in January 2010. Again, this discount is substantial relative to typical retail markup of 10 percent.

3.2 Estimates of Marginal Rates of Return from Bulk Discounts

We use the availability of bulk discounts to infer a lower bound on the average marginal rate of return. A histogram of purchase size (Figure 1) shows that shops do respond to the availability of bulk discounts by trying to make purchases that just exceed the discount thresholds -- there are bumps in the distribution at the cutoffs. However, a substantial fraction of purchases also fall in the intervals just below the discount thresholds, and shops frequently forgo the discounts they could achieve by buying a larger quantity of goods up front. On the other hand, shops sometimes do get discounts for total purchase amounts just below the threshold and we observe these discounts. Figure 2 presents the distribution of purchase sizes excluding data within 150 Ksh of threshold.

We can use this data to calculate a rough bound on the interest rates. If a shop does not buy

enough to obtain a discount, this implies that the cost of financing the incremental purchase is greater than the benefit from the discount. We calculate the rate of return that each shop could have realized had it bought goods earlier in order to obtain the bulk discount, given that it would have been able reproduce the same sales pattern going forward as the sales pattern that is empirically realized. If firms buy from wholesalers, we will not see this in our data. This will tend to downwardly bias our estimates since shops actually reorder more than we observe. In general, our analysis does not let us estimate the interest rates, but only allows us to bound the marginal rate of return – the advantage however, is that there is less of a concern about selection into the sample, since this is administrative data.

For example, suppose that a shop makes a 4,500 Ksh purchase each month. Given an interest rate of r over a period of a month, their cost of borrowing to get to 5,000 Ksh would be $500*r$. The benefit would be a discount of $0.005*5000$. If they are not borrowing to get to the 5,000 Ksh threshold, this implies that $500*r > 0.005*5000$, or $r > 0.05$. A 5 percent rate of return over one month would be equivalent to an annual rate of 82 percent.

Some shops have low turnover and buy very few goods from this distributor. For smaller shops it takes a long time to ever buy to the first threshold, so they are more likely to buy smaller quantities from wholesaler rather than distributor. Therefore, we cannot bound their rates of return very tightly. On average, most shops purchase from distributors about once every two weeks.

We restrict our sample to shops for which we observe at least 20 purchases in the data. Our estimates are therefore driven by those who purchase enough for us to bound their returns with some precision. These shops are generally larger than other shops, and have been in operation longer. To the extent that there are diminishing returns to scale, this sample will have lower underlying rates of return than the unrestricted sample, and the bounds we present should be regarded as lower bounds on the distribution of rates of return for the entire population of retailers.

To the extent that larger shops are likely to have been in operation longer, this sample will likely exclude new shops that may not yet have learned to take advantage of the discount, for which we might calculate a spuriously high bound for the rate of return due to a lack of information about the discounts.

We then search for the date on which they make a purchase that is closest to the next discount threshold. Using subsequent purchases, we then calculate the rate of return they could achieve by increasing the size of the purchase order to meet the next discount threshold. Using this method, we are able to bound rates of return for several thousand shops.

In order to get tight bounds, we look at the biggest unrealized return. However, to avoid being driven only by occasionally calculation errors, for instance rare "mistakes", we also look lower in the distribution.

We find evidence that rates of return for a significant fraction of the shops we study can be bounded at extremely high levels and observe high heterogeneity. Under the baseline perfect information assumption, a large percentage of shops have huge returns. Figure 4a shows the maximum observed (log) return for all shops. Since log returns are so massive for some shops, we focus on the bottom 80 percent of the distribution. Approximately 65% of shops have maximum returns of over 100% per year.

It could be that shops make rare errors, or only very occasionally face major liquidity constraints. If so, the maximum return might be substantially higher than it is in other periods. To address this, we calculate the (ex post realized) return to moving purchases forward, rank these from lowest to highest, and present the 95th percentile of this distribution. Results are presented in Figure 3, Panel B. The results show that over 35% of shops display returns over 100% on 1 out of every 20 purchases that they make.

To give a sense of magnitudes, a return of 100% per year can be interpreted as follows: at

median restocking frequency (14 days), this is roughly equivalent to purchasing 4,100 Ksh of products and restocking the other 900Ksh of inventory after 14 days. Similarly, a return of 50% per year is roughly equivalent to purchasing 3,600 Ksh and restocking the other 1,400 Ksh in 14 days. It is also striking that these results show vast heterogeneity in returns.

A major limitation of our analysis is that we calculate ex post realized returns. Thus we may pick up unexpected demand shops which we observe ex post but which were not anticipated by shops when they made their inventory choices. To make a rough adjustment for this, we assume shops thought it would take 50% or 100% longer to sell the items they bought (note that we are in the process of estimating a more sophisticated restocking model). Figure 4 shows how the estimated marginal returns under different adjustment scenarios. If we assume that it takes 50 to 100 percent longer to sell items that would have ordered, returns are still massive for many shops.

Another piece of evidence against unexpected demand shocks being the sole explanation is that, if it were, the massive heterogeneity in returns shown in the previous Figures should be driven largely by (unexpected) variance in demand. However, many of the shops with the highest returns are located in the same market centers as shops with far lower returns – in fact, regressing the bounded return on market fixed effects yields an R-squared close to 0. In addition, we can make a crude attempt at estimating the variance in returns by calculating total monthly purchases (which comprise 4 or 5 weekly visits from the distributor), and then calculating the shop-specific standard deviation in this measure. Regressing the bounds on the standard deviation gives a small, statistically significant coefficient, but the R-squared is again close to zero. Thus, while it may be the case that the same shops which fail to purchase the bulk discount also tend to face more heterogenous demand (perhaps directly because they are more often to stock out of certain products because their inventories are lower), it does not appear that exogenous demand fluctuations are driving the

relationship.[3] This same pattern appears to hold in our second approach (the stockout analysis), and shops directly report reasons other than demand fluctuations as the primary reason for keeping small inventories in a debriefing interview (Section 5).

There are several caveats to this analysis that should be noted. First, as can be seen from the figures, there is some bunching of purchases just below the discount threshold. In talking to the distributor, this appears to be due in part to the fact that discounts are given by truck drivers on their rounds, and at times in the past the distributor's system for checking that discounts were given correctly was not very effective. In particular, truck drivers seemed to have some discretion in giving discounts to shops that were just below the eligibility threshold. We are only able to address this by removing these purchases from the dataset (which, by potentially excluding some purchases just below the threshold for which the discount was not obtained, should bias our estimates downwards).

Second, we calculate very high bounds on rates of return at some point in time for these shops, but average rates of return across time may be substantially lower. Third, in the current version of the analysis, we account only roughly for uncertainty over which products will be in demand. Shops may delay purchasing products until some of the uncertainty becomes resolved. As we get product specific data in order, we plan to account for this by analyzing the expected returns to very simple investment rules -- for example, the return to increasing the purchase order by equal amounts for the three highest volume products. To the extent that shops have more information, and could have chosen a higher return bundle of goods to buy, this will be a lower bound on the rate of return they could have achieved by increasing the order to the next discount threshold.

[3] Note, however, that our dataset only includes purchases from the distributor. Shops may also restock from a wholesaler if stocks run low during the week, before the distributor returns. Thus our measure of total monthly sales is crude.

4. Estimating Marginal Rates of Return from Stockouts

4.1 Survey Data

Our second approach uses stock-outs of phone cards to estimate the marginal rates of return. The dukas for our sample were recruited from a census of small retailers in 16 small towns in Western Kenya. Shops were eligible to participate in the survey if they sold telephone cards, although a small number of businesses that sold these products but operate primarily as wholesalers were excluded from the sample. In addition, we excluded a small number of larger retail outlets (supermarkets") because they allow customers direct access to goods, so that the shopkeeper would have a difficult time observing and reporting the number of customers lost to stockouts. In total, 142 shops participated in the survey, amounting to 74% of shops in those areas. Our results can only be considered valid for the subset of shops that agreed to participate in the survey. However, as the focus of this section will ultimately be on heterogeneity in returns across firms, in order to demonstrate that rates of return are not equalized it is sufficient to show that the rate of return to a particular investment in a well-defined subset of firms differs from that in another set.

In total, the analysis to date includes data on 142 shops which were surveyed twice weekly for a period ranging from three months to one year. The survey collected information about the number of items sold that day, the last time the shop had restocked each item, and the number of customers who had been lost to stockouts for each product.

As noted above, we define the event in which a customer comes to ask for a product that is out of stock and does not purchase a substitute to be a "stockout". Daily data on stockouts for each item were constructed by asking shopkeepers to retrospectively report stockouts for each day since the previous survey. For some products, customers may substitute to another size or brand. To account for this, shopkeepers were asked whether the last customer on each day that requested a product that was out of stock substituted to another size or brand, or left. Substitution across brands

is extremely rare, as airtime can only be used on a phone with a sim card for that brand. Substituting to different sizes was more common, but still rare -- substitutions were reported in fewer than 6 percent of cases. In these cases, we conservatively do not count this event as a stockout (since a sale was still made, albeit to a different product).

Data on wholesale and retail prices for all goods were collected from the suppliers. Retail prices deviate somewhat from the prices reported by retailers for some products, but there are likely to be very few deviations in retail prices for phone cards, since the cards are printed with their value. Informal interviews with shopkeepers also indicate that deviations from the retail price are rare.

Since shops are visited by distributors at regular intervals, the relevant horizon over which shop owners decided how much inventory to hold is the interval between distributor visits. We thus aggregate the daily data to shop-product-distributor visit interval observations in order to impute the marginal rate of return. In order to construct the number of stockouts over one of these intervals, we must have data on the stockouts and on the exact dates of distributor visits. As a consequence, we drop observations that belong to an interval in which we cannot construct a complete history of stockouts; we also drop observations that fall in intervals of indeterminate length because the date of a past or future distributor visit is missing.

Table 2 displays summary statistics for the stockout sample. We observe each shop for a total of 263 days on average. Stockouts occur about 11 times per month for the average firm.

4.2 Empirical Methods

These data allow us to directly compute the expected rate of return to buying one incremental unit of inventory.

The net rate of return to holding an additional unit of inventory over the time between distributor visits can be expressed as:

$$r_i^I = \frac{(P^R - P^W)Pr(\omega > x_i) - (1 - \delta + c)}{P^W} \qquad (1)$$

where r_i^I is the marginal rate of return on the inventory investment for shop i over the distributor interval of I days, P^W and P^R are the wholesale and retail prices, respectively, ω is the number of customers who want to buy the product, x is the level of inventory, δ is the rate of depreciation and c is the cost of storage. The return to holding an additional unit is just the markup multiplied by the probability the marginal good sells less depreciation δ and the cost of storage c divided by the wholesale price. This calculation implicitly assumes that the firm values unsold cards at the wholesale price at the end of the period, less depreciation and storage costs.

Typically, it is difficult to measure rates of return, since the expected rate of return to an inventory increase depends not only on expected extra sales, but also on product depreciation, storage costs, the risk of theft, and the cross-elasticity of demand with respect to other products. For these reasons, the ideal product to study would be one for which depreciation, storage, and expected theft costs are minimal, and one which is neither a substitute nor complement for other goods sold by shops. For these reasons, we focus on top-up cards for pay-as-you-go cellular phone service which do not depreciate other than through inflation and are sufficiently small that the storage costs are negligible. Though theft is possible, no store in our survey reported any theft in the past year. Note that in these stores, phone cards and all other goods (with the possible exception of sodas) are typically kept behind the counter, so that customers do not have access to them unless they request the goods from the shopkeeper. Shops sell no substitutes for these goods other than top-up cards of other denominations, since the cards are specific to cell phone networks. They are unlikely to be strongly complementary to other household goods, but shops may incur some losses of sales of other products if they frequently stock out of phone cards due to a loss of reputation if customers prefer to buy all of their goods in one place. In this case, the estimates we present should be viewed

as lower bounds on the actual rate of return.

Empirically, another important advantage of such cards is that, since they are kept underneath the drawer, customers must ask the shopkeeper if has the product in stock (rather than being able to infer this from observing the shop). This will make it possible to estimate lost sales through a survey.

With this in mind, the analysis is straightforward. As there are minimal depreciation or storage costs, the last term in Equation (1) disappears. However, one small complication is that Equation (1) was for a single additional unit, whereas the minimum incremental order size is larger than this (usually 5 or 10 cards come in a pack). Taking into account the minimum order size, the marginal rate of return to holding an additional pack of cards over the period of the investment (the interval between distributor visits) in this context is then given by:

$$r_i^D = \frac{E[min(max(\omega - x, 0), N^{min}](P^R - P^W)}{P^W N^{min}} \qquad (2)$$

The only difference from the previous formula is that we have to account for the minimum order size N^{min}. The numerator reflects that now, by increasing inventory by N^{min} from x, the shopkeeper receives an additional markup (that he would not have otherwise) if demand is between x and $x + N^{min}$. If demand is greater than $x + N^{min}$, he receives the markup on only the N^{min} units (since requests beyond that point will result in stockouts). If demand is less than x, the shopkeeper receives no additional benefit. The minimum order size also affects the cost of investment, of course: the capital cost is now multiplied by the minimum order size in the denominator.

If the length of distributor visit intervals were constant across shops and across time, we could directly compute the expected marginal rate of return over those intervals from our data. In practice, the distributor visit intervals vary both within and across shops. For example, if a

distributor visits a shop on Tuesdays and Fridays every week, the data will consist of intervals of three days and intervals of four days.

We therefore construct daily interest rates. Since the return in Equations (1) and (2) is over the distributor interval of I days, the implied daily rate (r_i^d) is given by $r_i^I = (1 + r_i^d)^I - 1$. We then simply Taylor expand r_i^I around $r = 0$ to obtain an estimating equation that is linear in r:

$$r_i^I \approx ((1 + r_i^d)^I - 1)|_{r_i^d=0} + r_i^d I(1 + r_i^d)^{I-1} + H.O.T. \qquad (3)$$

$$\approx rI$$

We use this approximation estimate daily marginal rates of return for each shop using OLS regressions. To begin to interpret these estimates, we then estimate rates of return separately for each shop and phone card brand, since the standard theory would predict that rates of return should be equalized across products.

If cards are independent of other goods or if all shops face the same reputation costs from stockouts, then we can test for and estimate heterogeneity in marginal rates of return across shops. We do so by first using standard Wald tests of equality based on both the robust covariance matrix and the bootstrapped variance-covariance matrix. However, this test is not invariant to nonlinear transformations and since it relies on asymptotic approximations, may not be appropriate given the small number of shop-product-distributor visit intervals we observe for some shops in our data. Thus we also construct a nonparametric permutation test, in the spirit of a Fisher test, to check whether the observed distribution of estimands is consistent with what we would expect to observe in a world of equalized marginal rates of return. If marginal rates of return are equal across shops, there are no unobserved components of the marginal cost of holding an extra unit and no unobserved marginal benefits that vary across shops, and there is no autocorrelation in shocks to demand for a shop, then we can view the distribution of stockouts for all shops as the empirical distribution of residual shocks to demand for all shops.

Under these assumptions, we can generate distributions of the variation in estimated interest rates that would be realized if shops in fact faced the same interest rate and thus the same distribution of residual shocks to demand. We do this by randomly assigning shop-product-distributor intervals to artificial shops, and generating simulated distributions of estimated interest rates. We then compare the actual distribution of estimated interest rates to the simulated distributions. We generate simulated distributions of the variance of the estimated interest rates, the 90-10 spread, and the Wald test statistic test and compare the statistics for the actual distribution to the simulated distributions. We calculate the probability that the observed distribution of estimated coefficients would be generated at random under the null hypothesis of equal marginal rates of return by comparing the actual statistics (variance, 90-10 spread, and the Wald test statistic) to the empirical distributions of those generated by randomly permuting the shop assignment.

This procedure is robust to some types of correlation in shocks to demand over time. For example, if shocks to demand follow an AR(1) process and shops know this, then they will adjust their expectations accordingly. As a result, the *residual* shocks to demand for each shop will be uncorrelated over time.

Finally, we estimate the degree of underlying heterogeneity in rates of return in the population by using a random effects model. The object of interest is the standard deviation of estimated rates of return.

4.3 Results

The results of our empirical strategy is shown in Figure 5, which presents estimates of real rates of return and the standard errors (note that standard errors for the regressions are clustered at the shop level).. As there is significant heterogeneity in returns, Panel A plots the lower part of the distribution (below the 80^{th} percentile) while Panel B shows the implied rate of return for firms

above the 80th percentile. Overall, the average annualized real marginal rate of return is 33 percent, though this is heavily influenced by a few shocks with enormous returns - the median shop in the data faces an annualized real marginal rate of return close to 0. Only about 17.6 percent of shops have a return greater than 50%. Note that some rates of return are negative, since an estimated nominal rate of return of zero would imply a negative real rate of return (since our estimating equation takes into account the approximately 9% rate of inflation during this time period).

Why are these average returns so much lower than from the bulk analysis (and so much lower than other studies on the return to capital in developing countries)? A likely explanation is that there are significant costs to stocking out which are not captured in our estimating equation, in particular consumer goodwill. Our study period coincided with rapid growth in mobile phone usage in Kenya, and it may well be that stocking out of this product upset consumers.

However, while average returns are low, we use this method to principally examine heterogeneity. We find significant evidence that they are heterogeneous across shops. With a standard Wald test based on the robust covariance matrix or on the bootstrapped covariance matrix, we can reject the hypothesis that the estimated interest rates are equal across shops at the 1 percent level. Since the Wald test is not invariant to nonlinear transformations, we perform this test on the estimated coefficients, the daily rate of return, and the annual rate of return, and find similar results.

In addition, using the permutation test described above, we find that the 90-10 split of the observed distribution of estimated coefficients falls in the 99th percentile of the simulated distribution of 90-10 split (Figure 6). The actual 90-10 split of the estimated rates of return falls far above what would be expected if the shops actually faced the same rate of return. Taken together, we interpret these tests as a strong rejection of the hypothesis that the marginal rates of return to these investments are equal for the shops in our sample.

Given that we can reject homogeneity of returns across shops, we next estimate the extent

of the heterogeneity with a random effects OLS regression. Some of the variation in the distribution of fixed effects reflects sampling error, so we use a random effects model to estimate how much of this variation reflects real underlying heterogeneity in rates of return in the population of shops. We estimate that the standard deviation of the annual rate of return in the population is 98 percent.

The parametric assumption of normality in both the distribution of rates of return and the error term is almost certainly wrong, and the estimate of the extent of underlying variation changes dramatically with different assumptions about the distribution of random effects or with a random effects Poisson model. However, the conclusion that there is a large amount of underlying heterogeneity in the rates of return is qualitatively robust to the choice of specification -- the estimates of the underlying real variation in rates of return are consistently large. This provides evidence for economically significant departures from the equalization of rates of return across firms that would be predicted by the standard model.

As noted above, these results on heterogeneity should be interpreted carefully, as there may be unobserved heterogeneity in the costs of stockouts (such as lost sales of other goods or reputation costs) that could explain some fraction of the differences in rates of return across shops. Similar to the previous discussion, it could also be the case that there are underlying differences in the variance of residual demand across firms. Again, while there may be such variance, it seems unlikely to be exogenously imposed upon firms (since they are located in the same market centers with the same overall customer population). One test of this is that we strongly reject heterogeneity even when market fixed effects are included. Another is that the standard deviation of daily sales does not predict returns (and the R-squared, across firms, is close to zero).

We plan to test for across-firm differences in reputation costs by examining the cross-sectional relationship between the density of competitors in the immediate vicinity of the shop and the imputed rate of return, and also by using a difference-in-differences strategy to estimate the

impact of entry and exit of nearby competitors on stockouts. While not interpretable as causal estimates, these correlations would provide some idea of whether reputation costs are likely to be empirically significant in this context.

One additional issue in the current set of results is that our sample included only roughly 74% of the shops operating in the towns that we studied. However, our key results on both the level and variance are qualitatively robust to sample selection issues. In addition, rejecting the hypothesis of equal rates of return in the sample we do observe is sufficient to reject the hypothesis of equal rates of return in a larger sample.

5. Credit Constraints or Behavioral Anomalies?

Thus far, we have shown that returns are very high in our preferred approach, and heterogenous in both approaches. In this section, we provide some preliminary evidence on *why* the returns are as high as they are.

5.1. Correlation between background characteristics and returns

We first collected extensive data on the characteristics of shops and their owners, for shops in both samples. We were able to collect information on 113 shops in the stockout sample. However, as it was infeasible to survey the thousands of shops in the bulk analysis, we conducted surveys with approximately 480 which we could match to our administrative data (note that this sample is not selected other than for survivorship – we picked a subset of market centers in which to conduct surveys, and attempted to follow every firm in those centers). The shops were administered a detailed background survey which gathered information on a number of standard demographic and background measures such as the owner's age, sex, ethnicity, educational attainment, literacy, and the size of the owner's family. In addition, we collected exhaustive inventory data, documenting each

product and its value present in the shop at the time of the survey. The survey also included questions on the shopowner's access to savings and credit, land, durable good and other asset holdings, transfers he had given and received, other sources of income, self-reported credit constraints, and access to finance. Since trade credit provided by suppliers may also potentially be an important source of financing, a separate section of this questionnaire focused on the relationship between suppliers and the retailer, especially regarding any credit provided by suppliers.

In addition, we developed survey instruments to measure less commonly studied correlates of business success, including attitudes towards entrepreneurship, various psychological measures, and a host of measures of cognitive ability. This data provide valuable information on the entrepreneurial environment, how entrepreneurs think about investment decisions, and which characteristics best predict business success.

In order to estimate the relationship between owner characteristics and the annual rate of return, we ran univariate variance weighted least squares correlations, presented in Tables 4 and 5. Table 4 shows the correlates between estimated rates of return and firm owner traits and personality characteristics for the stockout analysis. Table 5 shows the same variables for the bulk discount analysis. Notably, these tables demonstrate that those background characteristics which would tend to be most correlated with credit constraints (such as education or asset ownership) do not correlate with the rates of return. This is a first piece of evidence that pure credit constraints seem unlikely to explain the whole story.

5.2. Correlation across products and methods

We next separately estimate rates of return implied by stockouts of Celtel and Safaricom products for each shop and compare these returns (for this analysis, we restrict attention to those shops which carry both brands). Figure 7 presents the results, which show that shops with very high

Safaricom returns have very high Celtel returns, and so there is a some correlation. But there is less correlation for lower range of returns, so the relationship between the two types of returns is not very strong. Formally, the hypothesis of equality of the two returns can be rejected (taking into account the sampling variation in both estimates).

A similar exercise can be conducted for the returns across the 2 methods (bulk discount and stockout). Though this exercise is complicated by the fact that shops are given different names in the two datasets (such that they are often hard to match directly), we again find evidence of a minimal correlation between the two methods.

5.3. Learning

We also look at time trends within our data to gather evidence to differentiate between the two competing hypotheses. To collect data on stockouts, we visited shops multiple times a week and measured lost sales. If shops were not optimizing, this data collection process could change behavior, by making the stock-out frequency more salient for the shop keeper. We thus look at how the rates of return varies with the amount of time that we have been following shops. Though the start time is not random, variation comes from the fact that market centers were rolled into the sample at difference times (primarily for convenience). Thus, the amount of time being followed is not perfectly collinear with overall time trends, and so we can control for both. Nevertheless, these results must be taken with great care given the non-random variation.

Table 6 shows the results. This regression is presented separately for the "old" sample (which contained very few firms and so was much less representative) and the "new" sample in which participation was higher. To get a shop-specific measure of learning, we regress the daily rate of return on the time trend since start of project (7/2005) and on time since each shop started stockout survey. The overall trend captures changes in the market. We find that returns decrease

over the duration of our study, strongly suggesting that the costs of stockouts had increased over time, likely due to increased risk of lost consumer goodwill to stocking out (the markup had not changed appreciably).

Interestingly, however, we observe significant downward time trends for the number of days since the first visit at a particular shop, even when controlling for these overall trends. This suggests that visits to administer the stockout survey indeed changed behavior, and that shopowners appear to have learned to stockout less frequently. The size of the learning effects can be observed intuitively, as follows: from column 1, reduce the annualized rate by about 1.5 point for every month in data. This is big at the mean but not very big for the 16% or so of shops at the right tail. There is also evidence for some time trends in the market, of roughly the same order of magnitude as the learning trend.

We examine these trends less parametrically by using the specific number of times a shop had been interviewed previously. To the extent that these visits serve as "reminders" that stockouts are important to keep track of, the visit count is the most direct way of testing for behavior change. We do this in Table 7 by regressing stockouts or the annualized daily return on the number of visits (with controls for the denomination and size of phone card and for the month and year of the interview). We find that the number of visits predicts fewer stockouts and a lower annualized return, though the effect does not begin to appear for a few visits.

Though potentially surprising, these results are similar to those found in an experiment among a range of small firms in Kenya, in which the length of time being followed was actually randomized. Beaman et al. (2011) surveyed firms about another related phenomenon – losing sales because of insufficient change – and find that reminding firms about this problem had a large effect on such behavior.

5.4. Debriefing Survey

To shed more light on the reasons for the observed behavior, we began a randomized controlled trial in 2010. In the trial, we randomly selected a subset of firms from the bulk discount analysis and performed an "information intervention" with them. In this intervention, we walked respondents through the bulk discount calculations. In addition, we collected some basic data from them on recent stockouts and on purchases from wholesalers (which are costly because their price is higher than the distribution firm), and used this to calculate rough back of the envelope estimates of the costs of carrying small inventories. We left it to the respondents to determine whether the cost exceeded this benefit. Since this treatment is randomized, we will be able to tell whether information alone had an effect by comparing post-treatment behavior of treatment and control firms.

Though we do not yet have results from the intervention, we do have some preliminary results from the intervention protocols and surveys themselves. The script to deliver the information material was extensively pre-tested to work in a participatory way, and to elicit feedback from respondents on their opinions of the returns. Though obviously quite speculative, the results are generally supportive of our other findings. These are presented in Table 8.

First, Panel A shows averages for the other costs of carrying small inventories (stockouts and having to purchase from a higher cost wholesaler). As can be seen, both are very common. Total lost profits are substantial though not enormous (weekly profits are around 5,000 Ksh per week, so that these costs account for around 1% of profits).

In Panel B, we asked respondents whether the returns were higher than expected. For all the costs, the vast majority report that the returns are higher than expected and that they are higher than the alternative investment opportunity.

We also asked respondents how they would come up with more if they wanted to stock more. Interestingly, most report that the money would come out of reduced consumption.

A final figure of note is in Panel C. At the time of this intervention, 84% percent of

respondents have accounts in a bank (and all of those with accounts use them). As the returns to these accounts are usually just a few percentage points a year, it appears that at least in some periods, respondents seem to be foregoing significant profits in the business to keep money in their savings account.

5.5. Firm Survival

We also look at firm survival in both samples and find some evidence to support the anecdotal findings that there is high turnover of the shops that are less successful. This suggests and some entrepreneurs enter the market, make errors and exit. Table 9 presents the data on firm survival correlated with estimated returns in the stockout sample. We find no such evidence in the bulk discount sample (though note that these survivorship results are over different lengths of time). We find some evidence that there is some competitive pressure at play, but evidently it is too weak to completely drive out those firms which experience the highest returns (and remove such massive heterogeneity).

6. Conclusion

We use evidence from inventories to provide a novel look at the marginal rates of return to investments available to rural retail enterprises in developing countries. With administrative data on a large sample of shops, we use bulk discounts to estimate a lower bound for the marginal annualized rate of return to capital for the median shop well over 100 percent per year. Our second (stock-out) method of estimation yields much lower average returns (likely due to unmeasured costs of stockouts, notably consumer goodwill), we find evidence for substantial heterogeneity in marginal rates of return among these shops -- using several tests, we reject the hypothesis that the estimated marginal rates of return are equal across shops.

This suggests the potential gains from improving the allocation of capital may be large. The ability to realize these gains and the policy levers most conducive to doing so depend on the sources of these differences. There are of course multiple potential hypotheses about why we observe such high and heterogeneous marginal returns to capital and why rates of return will not be equalized. The first explanation is the hypothesis that credit constraints prevent small shopkeepers from borrowing to equalize returns with the outside credit market. The second is the hypothesis that behavioral factors limit the ability of small entrepreneurs to equalize rates of return across different items within their firms. Our preliminary analysis suggests that credit constraints are not the only explanation: we find evidence for learning over time and do not find shop or owner characteristics to predict returns. We also do not find any evidence that the returns in the stockout analysis are correlated with self-reported credit constraints or with assets, suggesting that access to finance is not the explanation for our results. Our findings strongly suggest that shop owners are not fully optimizing.

Evidence that some shops do not optimize is related to broader issues of development. In developed economies, firms make mistakes too, but efficient equilibrium is arguably quickly achieved as efficient producers expand while inefficient producers are replaced. Though we do find evidence that firms with the highest returns do exit in Kenya as well, the effect is not strong. It maybe that the market mechanism may not work as well in developing contexts, such as Kenya, due to labor monitoring problems; government expropriation risk; and retail price maintenance, including manufacturer facilitated price collusion to keep markup high, induce retailers to carry brand, where free entry can lead to excess entry.

We measure the rate of return to investment in a narrow category of activities, and this is sufficient to reject the standard model. Under stronger assumptions, the rate of return we measure also provides information about the rate of return to a broader set of investments. The marginal rate

of return we measure may also reflect the marginal rate of return to capital in a broad swath of rural economic activities if the individuals we study (or the households to which they belong) are diversified and allocate their working capital across a set of productive activities (such as farming, raising poultry, etc). Diversification has important implications for the interpretation of the estimand not only for this reason, but also because if these shopowners are diversified, it may be possible to interpret this rate of return as the social marginal rate of return rather than just the private rate of return. Aggregate stockouts in these market towns are rare and there are typically many shops selling the same goods, so in the context of rural retail shops, the social return to financing the purchase of an additional unit of inventory by any one shop may be close to zero -- if a customer finds that one shop has stocked out of a particular product, he will buy from a competitor. However, if shop owners are diversified and participate in a variety of productive activities including some that do not exhibit this zero-sum feature, then the marginal rate of return we measure may reflect the social marginal rate of return to capital as well as the private return.

6. Bibliography

Aleem, I. (1990): "Imperfect Information, Screening and the Costs of Informal Lending: A Study of a Rural Credit Market in Pakistan," World Bank Economic Review, 3, 329–349.

Anagol, S., and C. Udry (2006): "The Return to Capital in Ghana," *American Economic Review Papers and Proceedings*, 96(2), 388.

Banerjee, A., and E. Duflo (2008): "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," Mimeo, MIT.

Banerjee, A., and K. Munshi (2004): "How Efficiently is Capital Allocated: Evidence from the Knitted Garment Industry in Tirupur," *Review of Economic Studies* 71 (1) 19-42.

Beaman, Lori, Jeremy Magruder, and Jonathan Robinson (2011), "Small change and forgotten sales: Evidence from an experiment with small businesses in Kenya," working paper.

Caselli, F., and J. Feyrer (2007): "The Marginal Product of Capital," *Quarterly Journal of Economics* 122 (2): 535-568

Daniels, L., and D. Mead (1998): "The Contribution of Small Enterprises to Household and National Income in Kenya," Economic Development and Cultural Change, 47(1), 45.

de Mel, Suresh, David McKenzie and Christopher Woodruff (2008), Returns to Capital in Microenterprises: Evidence from a Field Experiment, *Quarterly Journal of Economics* 123 (4): 1329-1372.

Fafchamps, Marcel, David McKenzie, Simon R. Quinn, and Christopher Woodruff (2011). "When is capital enough to get female microenterprises growing? Evidence from a randomized experiment in Ghana." NBER Working Paper No. 17207.

Hsieh, Chang-Tai and Peter Klenow (2009), "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics* 124 (4): 1403-1448.

Lucas, R. E. (1990): "Why Doesn't Capital Flow from Rich to Poor Countries?," *American Economic Review Papers and Proceedings*, 80(2), 92.

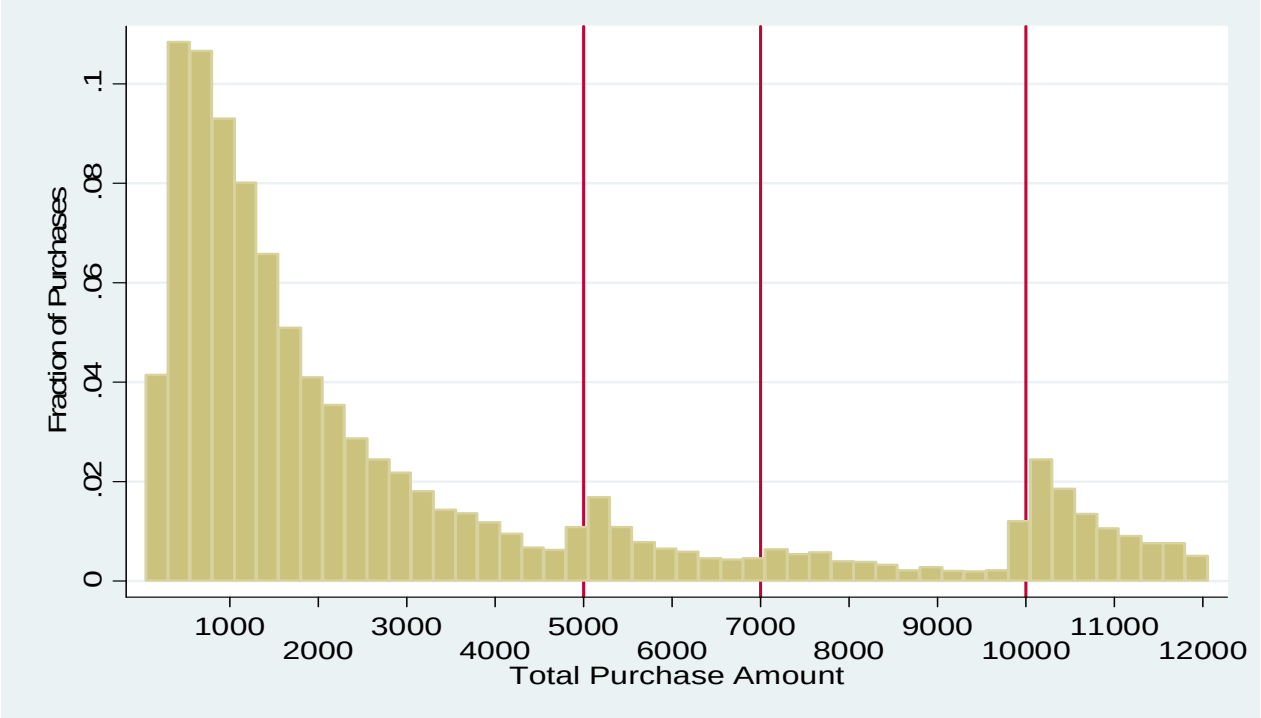
McKenzie, D., and C. Woodruff (2006): "Do Entry Costs Provide an Empirical Basis for Poverty Traps? Evidence from Mexican Microenterprises," *Economic Development and Cultural Change* 55 (1): 3-42.

McKenzie, David and Christopher Woodruff (2008), "Experimental Evidence on Returns to Capital and Access to Finance in Mexico." *World Bank Economic Review* 22 (3): 457-482.

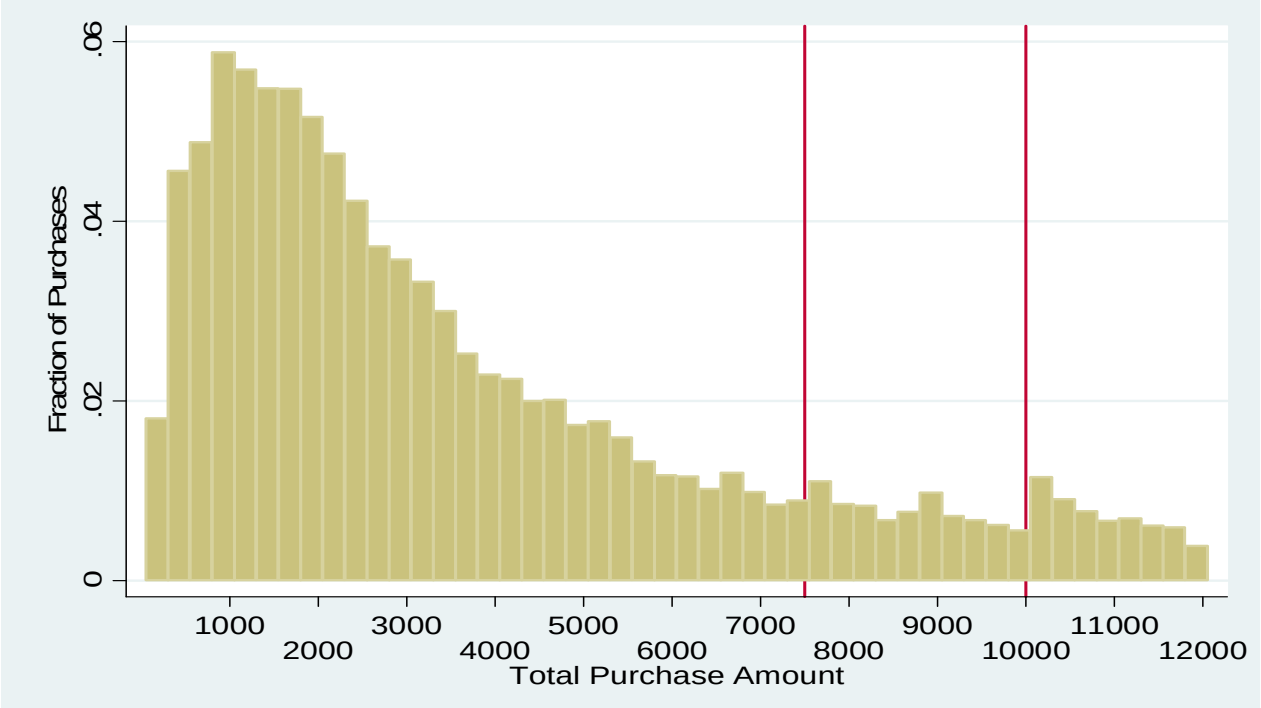
Zinman, J. (2002): "Do Credit Market Interventions Work? Evidence from the Community Reinvestment Act," Mimeo, Federal Reserve Bank of New York.

Figure 1. Purchase Quantities, Distributor Data

Panel A. Before 2010



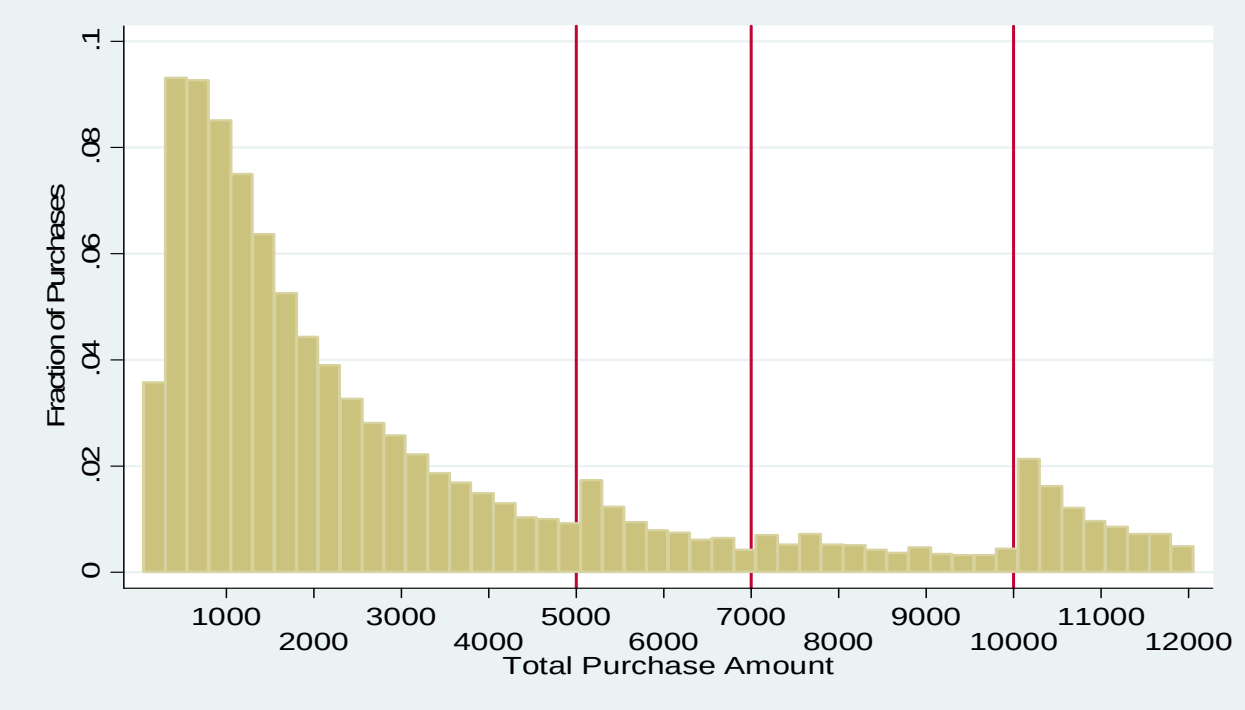
Panel B. After 2010



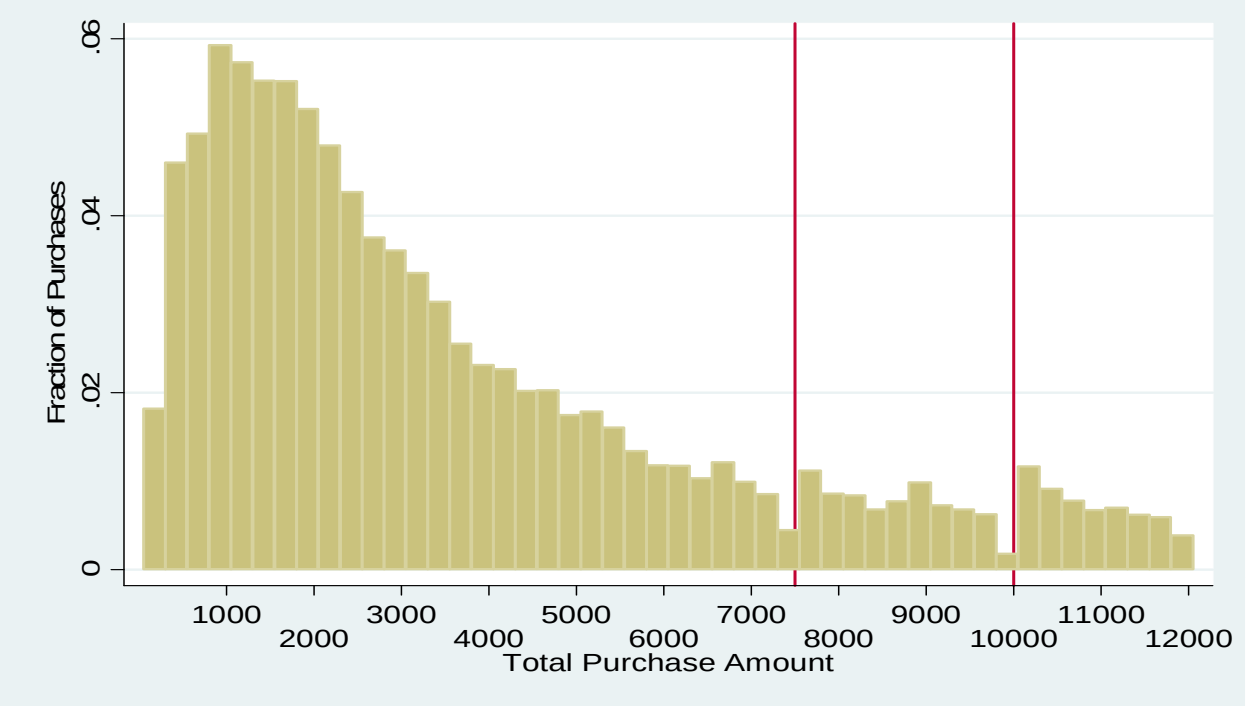
Notes: From distributor database. The discount schedule was 0.5, 1, and 1.5% off at 5,000, 7,0000, and 10,000 Ksh prior to 2010. After 2010, the thresholds changed to 7,500, 10,000, and 20,000 Ksh. Figure shows purchases below 12,000 Ksh.

Figure 2. Purchase Quantities Excluding Mismeasured Quantities, Distributor Data

Panel A. Before 2010



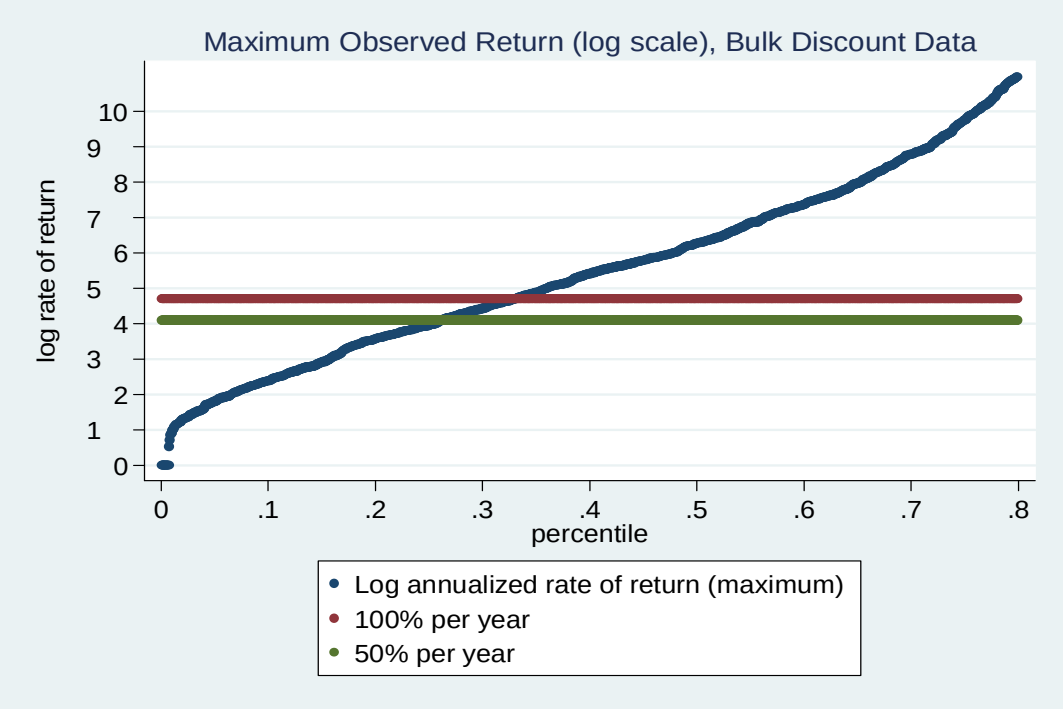
Panel B. After 2010



Notes: From distributor database. Purchases within 150 Ksh of the threshold are removed. The discount schedule was 0.5, 1, and 1.5% off at 5,000, 7,0000, and 10,000 Ksh prior to 2010. After 2010, the thresholds changed to 7,500, 10,000, and 20,000 Ksh. Figure shows purchases below 12,000 Ksh.

Figure 3. Estimated Returns, Bulk Discount Data

Panel A. Maximum Observed Return



Panel B. 95th Percentile of Returns

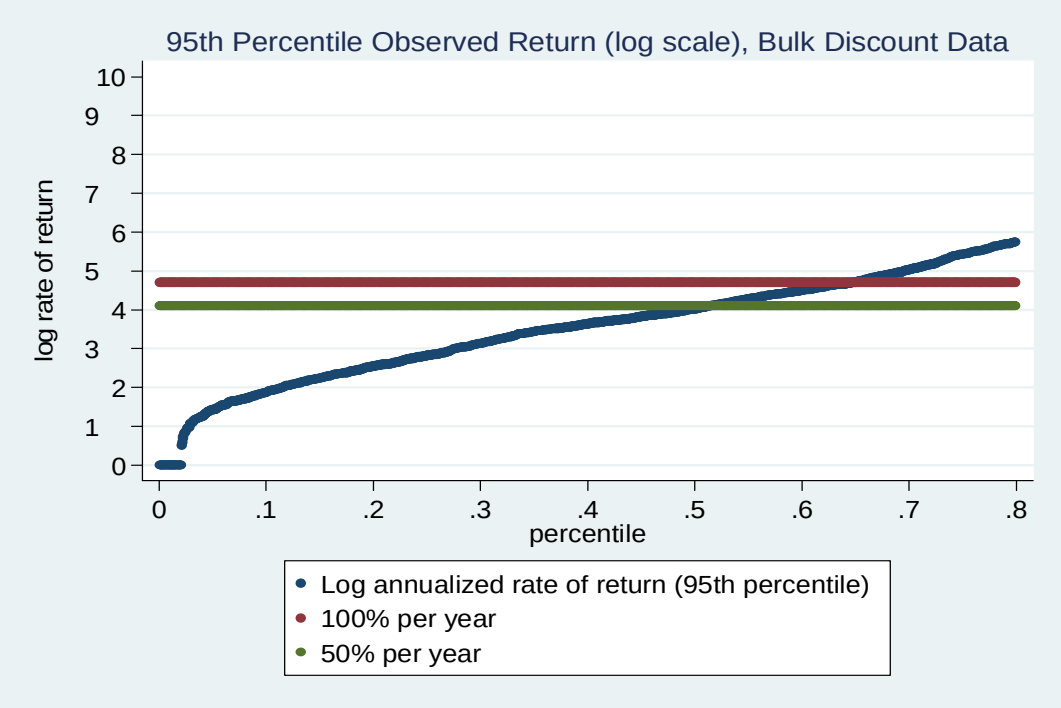
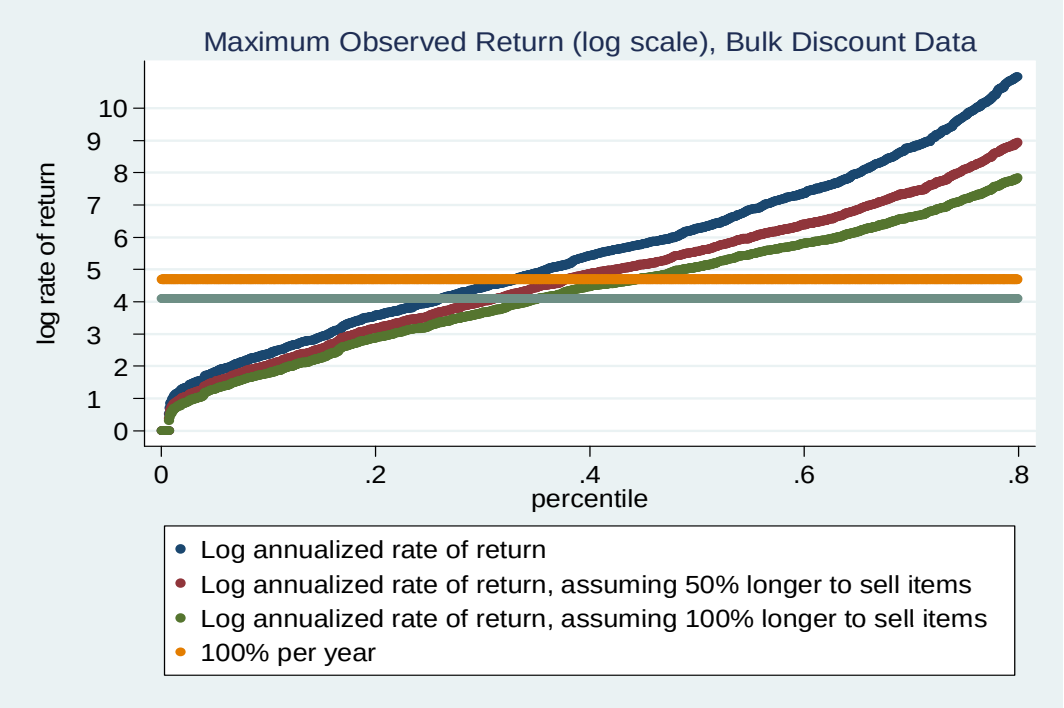


Figure 4. Rough Adjustment for Unexpected Demand Shocks

Panel A. Maximum Observed Return



Panel B. 95th Percentile of Returns

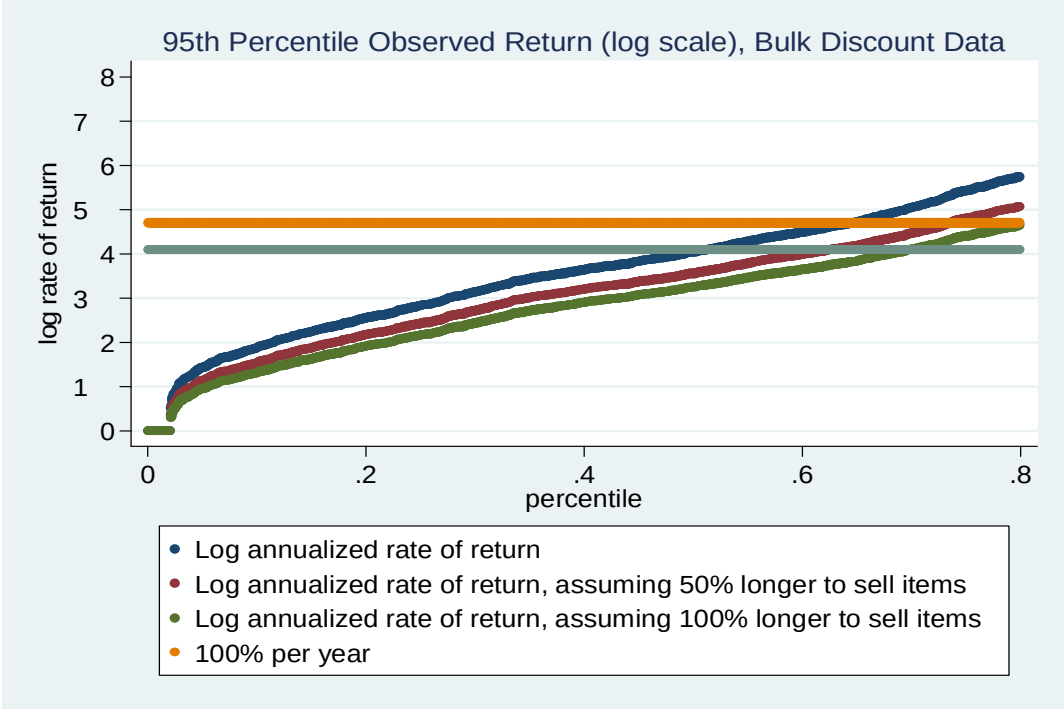
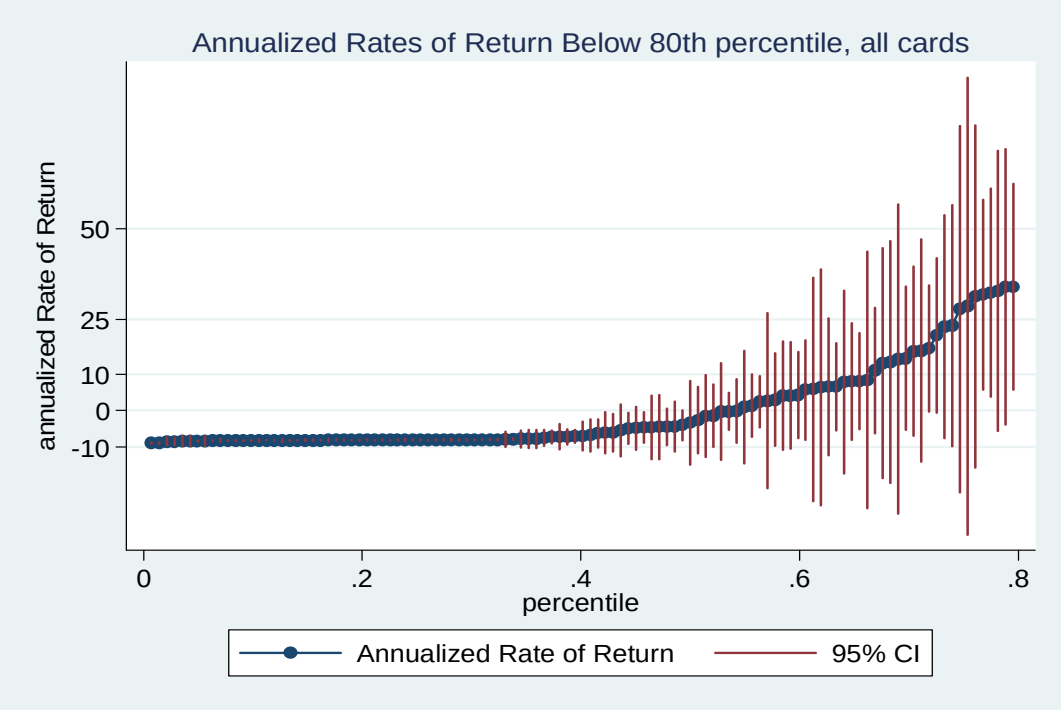


Figure 5. Estimated Annualized Returns (Stockout Data)

Panel A. Below 80th percentile



Panel B. Above 80th percentile

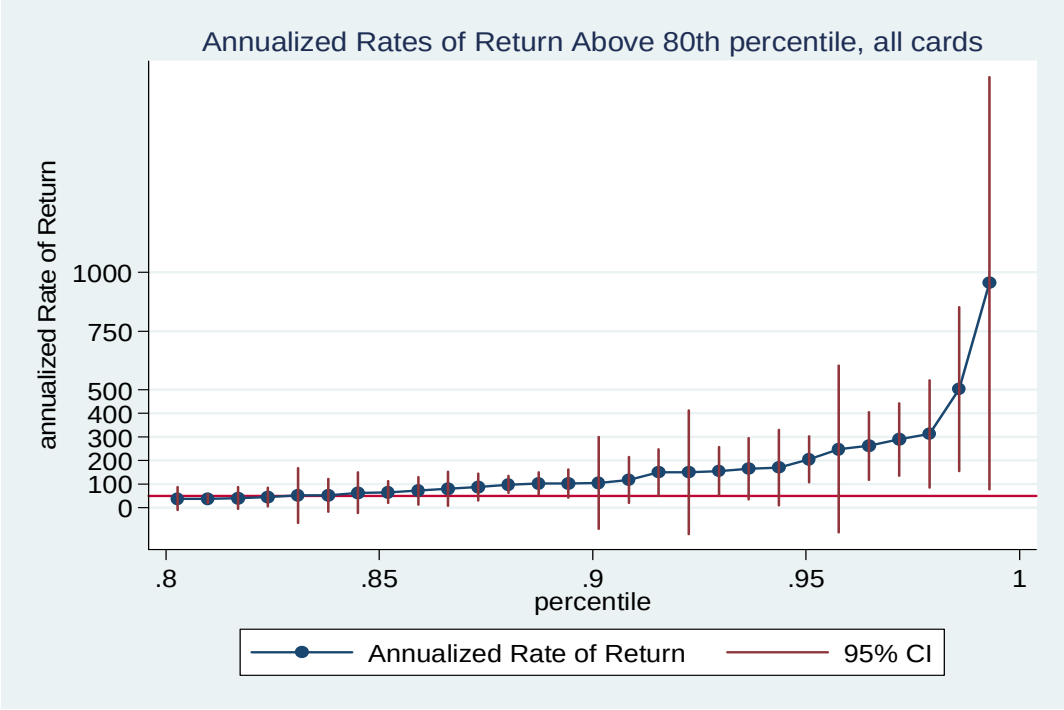
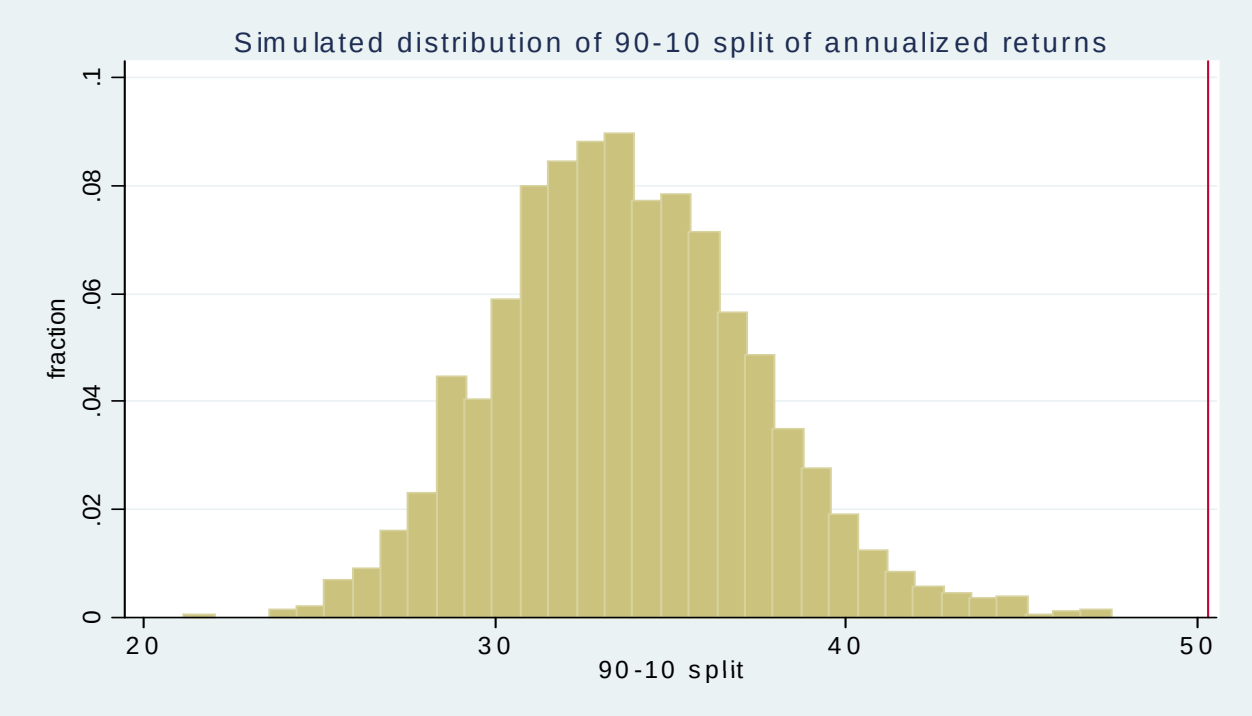
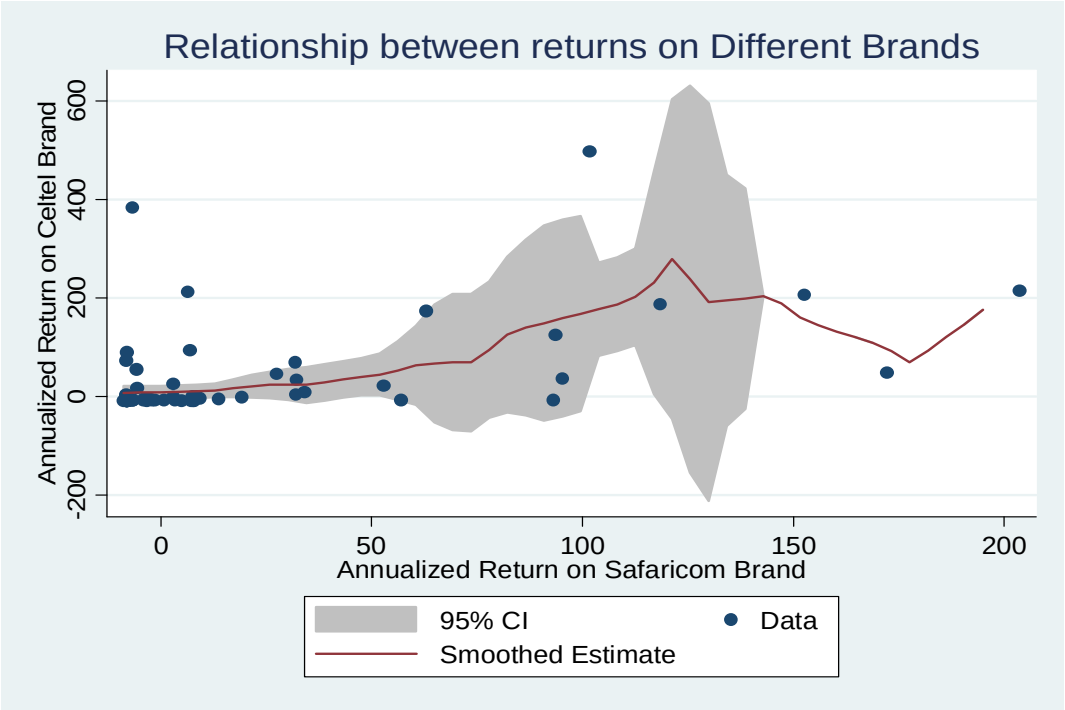


Figure 6. Permutation Test (Stockout Data)



Notes: Red line represents actual 90-10 split.

Figure 7. Relationship Between Returns on Different Brands (Stockout Data)



Notes: Graph includes only those retailers with estimated returns less than 500 percent per year.

Table 1. Summary Statistics, Distributor Data

	(1) Mean
Number of Purchases Observed	63.02
	(39.36)
Purchases per month	2.19
(note: truck visits once per week)	(0.89)
Value of Average Purchase	3731.00
	(9176.00)
Purchase Frequency (days)	17.03
	(9.09)
Total Days in sample	916.46
	(448.73)
Number of Firms	2136

Table 2. Summary Statistics, Stockout Data

	(1) Mean	(2) N
Number of Intervals per Shop	72.15	142
	(53.15)	
Number of Days	262.64	142
	(190.35)	
Number of Denomations of Phone Cards	3.51	142
	(1.57)	
Stockouts per month	10.69	142
	(23.73)	
Number of days per distributor visit	32.08	121
	(15.06)	
Probability of Restocking during visit	3.64	10245
	(1.65)	

Table 3. Background Characteristics, Stockout Analysis

	(1) Mean	(2) N
Years of Education	10.92	113
	(3.22)	
Years Shop Open	6.46	111
	(6.48)	
Male	0.48	113
	(0.50)	
Age	34.19	112
	(9.65)	
Can Read and Write English	0.96	113
	(0.21)	
Has Bank Account	0.66	113
	(0.47)	
Owns Land	0.88	113
	(0.33)	
If yes, acres of Land Owned	4.54	95
Participates in ROSCA	0.45	112
	(0.50)	
If yes, contributions in past year	95.01	112
	(164.35)	
Got loan from bank or microfinance institution in past year	0.28	113
	(0.45)	
If yes, value of loans	2198.44	32
	(1983.52)	
Sent transfer to friend / family	0.67	113
	(0.47)	
If yes, amount	143.92	77
	(206.46)	
Got transfer from friend / family	0.65	96
	(0.48)	
Wants more Credit but Can't Get it	0.76	112
	(0.43)	
Got phone cards on Credit in Past Yr	0.05	113
	(0.23)	
Got household goods on credit in past Yr	0.06	113
	(0.24)	
Digits Forward Score	7.33	106
	(1.96)	
Amount Invested in Risky Asset	65.05	106
	(17.86)	
Time consistent	0.60	91
	(0.49)	
Present-biased	0.08	91
	(0.27)	
Patient now, impatient later	0.22	91
	(0.42)	
Completely impatient in present and future	0.10	91
	(0.30)	
Think about business all of the time	0.22	107
	(0.42)	
Commonly miss family events for work	0.23	107
	(0.43)	
Very willing to sell assets to keep business going	0.07	107
	(0.25)	
Never keeps financial books	0.29	113
	(0.29)	
Estimated weekly revenue (USD)	600.92	108
	(1537.65)	
Estimated weekly investment (USD)	209.92	109
	(1406.07)	
Time to finish mazes (seconds)	148.83	115
	(128.67)	
Raven's matrix (percent)	0.73	115
	(0.16)	
English score (percent)	0.53	114
	(0.19)	
Math score (percent)	0.52	115
	(0.17)	

Table 4. Correlates of Returns, stockout analysis

	(1) Depvar: log(estimat	(2) N
Years of Education	-0.03	113
	(0.05)	
Years Shop Open	-0.02	111
	(0.03)	
Male	-0.13	113
	(0.31)	
Age	-0.02	112
	(0.02)	
Can Read and Write English	-0.08	113
	(0.52)	
Has Bank Account	0.01	113
	(0.33)	
Owns Land	-0.49	113
	(0.54)	
If yes, acres of Land Owned	0.02	95
	(0.05)	
Participates in ROSCA	-0.24	112
	(0.31)	
If yes, contributions in past year	0.00	112
	0.00	
Got loan from bank or microfinance	0.02	113
institution in past year	(0.34)	
If yes, value of loans	0.00	32
	0.00	
Sent transfer to friend / family	0.01	113
	(0.33)	
If yes, amount	-2.06	77
	(0.89)**	
Got transfer from friend / family	0.20	96
	(0.35)	
Wants more Credit but Can't	0.26	112
Get it	(0.34)	
Got phone cards on Credit in	0.48	113
Past Yr	(0.80)	
Got household goods on credit	-0.08	113
in past Yr	(0.57)	
Digits Forward Score	0.24	106
	(0.17)	
Amount Invested in Risky Asset	0.01	106
	(0.01)	
Think about business all of the time	-0.72	107
	(0.31)**	
Commonly miss family events for	-0.05	107
work	(0.37)	
Very willing to sell assets to	-0.48	107
keep business going	(0.41)	
Never keeps financial books	0.15	113
	(0.32)	
Estimated weekly revenue	-0.15	108
	(0.05)***	
Estimated weekly investment	-0.10	109
	(0.02)***	
Time to finish mazes (seconds)	-0.19	115
	(0.15)	
Raven's matrix (percent)	0.26	112
	(0.34)	
English score (percent)	-0.05	114
	(0.16)	
Math score (percent)	0.16	115
	(0.18)	

Notes: Each row reports a univariate regression of returns on the given variable. Standard errors in parentheses, *, **, *** indicate significance at 10, 5, and 1% respectively.

Table 5. Correlates of Returns (Bulk Discount Analysis)

	(1)	(2)
	pvar: log 10+estimated rate of return (m	
	Estimate	Obs
Years of Education	0.16	466
	(0.10)	
Years Shop Open	0.05	466
	(0.06)	
Male	-1.57	468
	(0.661)**	
Age	-0.03	466
	(0.04)	
Can Read and Write English	1.63	446
	(4.20)	
Has Bank Account	2.47	439
	(0.924)***	
Owns Land	0.47	468
	(0.77)	
If yes, acres of Land Owned	0.06	355
	(0.13)	
Participates in ROSCA	-0.72	466
	(0.68)	
If yes, contributions in past year	0.00	206
	0.00	
Got loan from bank or microfinance	-1.57	468
institution in past year	(0.661)**	
If yes, value of loans		
Sent transfer to friend / family	-0.38	430
	(0.81)	
If yes, amount		
Got transfer from friend / family	1.40	455
	(0.675)**	
Wants more Credit but Can't	-1.76	468
Get it	(0.672)***	
Digits Forward Score	-0.08	470
	(0.18)	
Think about business all of the time	-0.17	472
	(0.76)	
Commonly miss family events for	0.04	472
work	(0.73)	
Never keeps financial records	-0.38	468
	(0.82)	
Estimated weekly revenue (1000 Ksh)	0.01	413
	(0.002)***	
Estimated weekly investment (1000 Ksh)	0.04	466
	(0.11)	
Time to finish mazes (seconds)	0.00	481
	(0.00)	
Raven's matrix (percent)	0.66	484
	(1.85)	
English score (percent)	0.20	474
	(1.06)	
Math score (percent)	3.30	484
	(1.861)*	

Notes: Each row reports a univariate regression of returns on the given variable. Standard errors in parentheses. *, **, *** indicate significance at 10, 5, and 1% respectively.

Table 6. Time Trends in Estimated Returns

	(1)	(2)	(3)	(4)	(5)	(6)
Days Since first Visit at Shop	-0.049	-0.050	-0.037	-0.096	-0.011	-0.096
	(0.006)***	(0.006)***	(0.004)***	(0.024)***	(0.008)	(0.024)***
Days since Project Start (July, 2005)	-0.047	-0.05	-0.05	-0.172	-0.013	-0.172
	(0.007)***	(0.007)***	(0.013)***	(0.029)***	(0.009)	(0.029)***
Days Since first Visit at Shop						0.059
* New Sample						(0.024)**
Days since Project Start (July, 2005)						0.123
* New Sample						(0.032)***
Dummy for New Sample						-32.246
						(11.947)**
Controls for Card Denomination/Brand	N	Y	N	N	N	N
Observations	9503	9503	7959	1544	1529	9503
Mean of Dependent Variable	10.61	10.61	7.98	25.11	0.49	10.61
Sample	Both	Both	New	Old	Only those in both new & old samples	Both

Notes: report (annualized) regression coefficients from regression of daily interest rate on these independent variables. Std. errors calculated by delta method.
p-value for test that Days Since first Visit at Shop + Days Since First Visit at Shop * New Sample <0.001

Table 7. Stockouts and Estimated Returns as Function of Number of Visits

	(1)	(2)	(3)	(4)
Panel A. Stockouts				
Number of Visits	-0.025	-0.206	-0.026	
	(0.017)	(0.071)***	(0.017)	
Dummy for Visit = 2				0.008
				(0.418)
Dummy for Visit = 3				-0.213
				(0.329)
Dummy for Visit = 4				-0.507
				(0.339)
Dummy for Visit = 5				-0.757
				(0.354)**
Visits Included	All	First 5 Visits	First 5	First 5
Observations	9467	1103	8701	1103
Number of shops[1]	142	121	121	121
Panel B. Estimated Annualized Return				
Number of Visits	-0.014	-0.155	-0.014	
	(0.012)	(0.053)***	(0.012)	
Dummy for Visit = 2				0.228
				(0.283)
Dummy for Visit = 3				-0.007
				(0.284)
Dummy for Visit = 4				-0.318
				(0.224)
Dummy for Visit = 5				-0.470
				(0.229)**
Visits Included	All	First 5 Visits	First 5	First 5
Observations	9467	1103	8701	1103
Number of shops[1]	142	121	121	121

Notes: Standard errors, clustered at the individual level, in parentheses. Regressions include time and product denomination controls. *, **, and *** indicate significance at 10, 5, and 1% respectively.

[1]There are fewer shops in the first 5 visits because information was filled incompletely for some shops early in the data collection period.

Table 8. Debriefing survey (after information intervention)

	(1) Mean	(2) N
Panel A. Other Costs of Holding Small Inventories		
Do you ever stock out?	0.90	212
Lost profits from stockout in past week	326.49	211
	(326.49)	
Do you ever purchase from a wholesaler?	0.35	210
Lost profits from buying wholesale in past month	89.35	212
	(89.35)	
Panel B. Post Information Intervention		
Would you say it's common that you buy close to the bulk threshold without reaching it?	0.79	82
Are losses more than expected?		
Stockouts	0.77	210
Purchasing from Wholesaler	0.75	71
Failing to qualify for bulk discount	0.53	83
Are returns from buying more inventory higher than alternative?		
Stockouts	0.84	212
Purchasing from Wholesaler	0.90	68
Failing to qualify for bulk discount	0.89	82
The last time you didn't buy up to the threshold, why?		
Hadn't realized how high the returns were	0.66	82
Another investment has higher returns	0.15	82
Emergency came up	0.26	82
Demand is unpredictable	0.10	82
If you were to stock more, where would money come from?		
Invest less in other product	0.12	212
Reduce consumption	0.71	212
Other business	0.21	212
Some other investment	0.15	212
Other	0.13	212
Panel C. Other Investments		
Do you have a savings account in a bank?	0.84	212
If yes, do you have money in it?	0.99	178

Table 9. Firm Survival as Function of Returns

	(1)	(2)	(3)
Panel A. Stockout Sample	Dependent Variable: Firm still Open in April 2010		
Log (Annualized Rate of Return + 10)	-0.05		
	(0.025)**		
Dummy for Annualized Rate of Return > 30		-0.187	
		(0.104)*	
Dummy for Annualized Rate of Return > 50			-0.225
			(0.118)*
Constant	0.789	0.706	0.699
	(0.074)***	(0.051)***	(0.049)***
Observations	112	112	112
Mean of Dependent Variable	0.68	0.68	0.68
	(1)	(2)	(3)
Panel B. Bulk Discount Analysis	Dependent Variable: Firm still Open in May 2011		
Maximum Return			
Log (Annualized Rate of Return + 10)	0.004		
	(0.002)**		
Dummy for Annualized Rate of Return > 50		0.060	
		(0.033)*	
Dummy for Annualized Rate of Return > 100			0.085
			(0.030)***
	0.881	0.871	0.871
	(0.021)***	(0.026)***	(0.026)***
Observations	459	459	459
Mean of Dependent Variable	0.92	0.92	0.92
Mean of Independent Variable	9.59	0.82	0.76
Maximum Return			
Log (Annualized Rate of Return + 10)	0.008		
	(0.01)		
Dummy for Annualized Rate of Return > 50		0.014	
		(0.027)	
Dummy for Annualized Rate of Return > 100			0.027
			(0.025)
Constant	0.892	0.881	0.881
	(0.036)***	(0.021)***	(0.021)***
Observations	414	414	414
Mean of Dependent Variable	0.93	0.93	0.93
Mean of Independent Variable	0.76	0.66	0.48

Notes: In Panel A, sample consists of shops identified in 2007. In Panel B, sample consists of firms that were interviewed in 2010.